

PCT

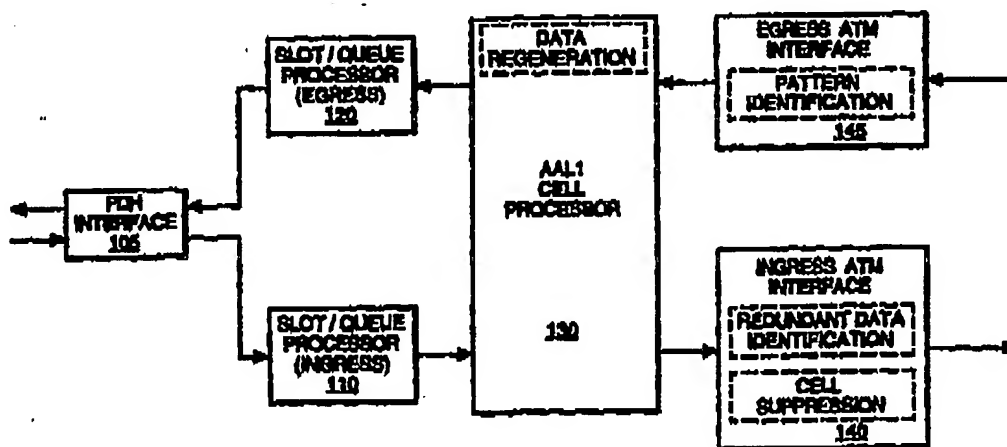
WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6: H04Q 11/00		A2	(11) International Publication Number: WO 98/51113
			(43) International Publication Date: 12 November 1998 (12.11.98)
(21) International Application Number: PCT/GB98/01243		(51) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, GW, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).	
(22) International Filing Date: 29 April 1998 (29.04.98)			
(30) Priority Data: 9709110.2 2 May 1997 (02.05.97) GB 08/896,380 18 July 1997 (18.07.97) US			
(71) Applicant (for all designated States except US): NORTHERN TELECOM LIMITED [CA/CA]: World Trade Center of Montreal, 8th floor, 380 St. Antoine Street West, Montreal, Quebec H2Y 3Y4 (CA).			
(72) Inventors; and (75) Inventors/Applicants (for US only): DOLBY, Rudi, Benjamin [GB/GB]; 55 Greenhill Park, Bishop's Stortford, Hertfordshire CM23 4EW (GB). HARVEY, Giles [GB/GB]; 201 Foldcroft, Harlow, Essex CM20 1SW (GB). JENKINS, Nigel, Phillip [GB/GB]; 15 Holmes Meadow, Harlow, Essex CM19 5SG (GB). RAVIRAJ, Rajakulasingam [LK/GB]; 54 Sharpcroft, Harlow, Essex CM19 4AB (GB).			
(74) Agent: RYAN, John, Peter, William; Nortel Patents, London Road, Harlow, Essex CM17 9NA (GB).		Published Without international search report and to be republished upon receipt of that report.	

(54) Title: DATA SUPPRESSION AND REGENERATION



(57) Abstract

A method of reducing bandwidth used on a telecommunications link comprising, at an input to the link, the steps of: receiving a plurality of data packets, each packet comprising data and a packet identifier; determining packets which contain redundant data; transmitting, across the link, packets which do not contain redundant data; and further comprising, at an output of the telecommunications link, the steps of: receiving the transmitted packets; determining missing packets according to the identifiers of received packets; and generating data for the missing packets, which generated data corresponds to redundant data not transmitted across the link. This ensures bit count integrity between the input and the output of the link. Redundant data may be data following a predictable pattern, such as identical data. Where the packet identifiers are a recurring sequence of N symbols, the input transmits at least one packet in every N packets.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

DATA SUPPRESSION AND REGENERATION

TECHNICAL FIELD

This invention relates to a method of reducing bandwidth used on a telecommunications link, a method of operating an output of a telecommunications link on which bandwidth is being reduced, and to apparatus for performing these methods.

BACKGROUND OF THE INVENTION

Packet, cell and frame-based transport protocols are increasingly being used in telecommunications networks. In particular, Asynchronous Transfer Mode (ATM) technology is a flexible form of transmission which allows various types of service traffic, such as voice, video or data, to be multiplexed together on to a common means of transmission, the traffic being carried in cells each having a payload and a header indicating the destination of the cell. ATM statistically multiplexes traffic and removes the rigid relationship between traffic and the time-domain which exists in other multiplexing protocols. ATM can accommodate both traffic whose information rate is constant during a call, and traffic whose information rate varies during a call. Operators are investing in ATM transmission equipment to accommodate users' increasingly diverse communications requirements.

In order to minimise costs, a network operator wants to use his transmission network most efficiently. One way of achieving efficiency is to make the best use of the bandwidth of the network.

One known way of reducing the amount of bandwidth which a telephone call uses on a transmission network is to suppress the periods of silence which occur during a telephone conversation. Silence suppression detects silence at an input to a transmission network, does not send data representing silence across the network, and at an output of the network inserts 'comfort noise' in place of the removed silence. The bandwidth that is saved by not transmitting the silent periods of a call allows other traffic such as bursty data traffic to be transported across the network. This sharing of the network is particularly easy to achieve in a cell-based network.

A technique for saving bandwidth on a cell-based transmission network is used in Northern Telecom's PASSPORT™ switch. The PASSPORT™

switch looks for consecutive identical cells representing an idle channel condition. When enough identical cells are received, the switch stops transmitting cells onto the network. At the destination the PASSPORT™ switch continues to transmit the last cell that it received, thereby saving bandwidth on the network. In this scheme the data which is suppressed at the input to the network is not regenerated accurately at the output of the network.

SUMMARY OF THE INVENTION

The present invention seeks to provide an alternative method of saving bandwidth.

According to one aspect of the present invention, there is provided a method of reducing bandwidth used on a telecommunications link comprising, at an input to the link, the steps of:

- receiving a plurality of data packets, each packet comprising data and a packet identifier;
- determining packets which contain redundant data;
- transmitting, across the link, packets which do not contain redundant data;

and further comprising, at an output of the telecommunications link, the steps of:

- receiving the transmitted packets;
- determining missing packets according to the identifiers of received packets; and,
- generating data for the missing packets, which generated data corresponds to redundant data not transmitted across the link.

By not transmitting redundant data the packets use a reduced bandwidth of the telecommunications link.

The step of determining missing packets according to their identifiers allows redundant data to be regenerated more accurately at the output of the link. The output of the link is able to determine how many packets were suppressed at the input to the link, thus maintaining bit count integrity between the data arriving at the input of the link and the data flowing out of the output of the link.

5 The regenerated data may be placed in packets for onward transmission across a further telecommunications link. Alternatively, it may be simply be regenerated, such as where the regeneration is performed at the interface between a packet-based network and a non-packet network.

10 Preferably the step of determining packets which contain redundant data comprises identifying a sequence of consecutive packets containing data that follows a predictable pattern and transmitting only the first packet or packets in the sequence. In many cases it should be possible to transmit only the first packet in the identified sequence, thereby maximising bandwidth savings, but in some cases it is necessary to transmit the first few packets in the identified sequence so that the output can correctly identify that data needs to be regenerated. In its simplest form the input identifies a sequence of packets containing the same data. Each packet may comprise a plurality of identical units, such as an idle data flag, an alarm flag, or a group of data bytes. Rotated versions of the flags may be identified, the rotation being caused by the flags not being aligned exactly with the bytes of the packet payload. Another pattern which may be identified is a ramp in data values.

20 In situations where only a portion of each packet contains user data, it is advantageous that the step of determining packets which contain redundant data examines only the portion of the packet containing user data. This allows the technique to work with partially-filled cells, such as are used in delay-sensitive applications.

30 Preferably the packet identifiers of a sequence of packets are a sequence of symbols and the step of determining missing packets according to the identifiers of received packets comprises identifying missing symbols in the sequence of identifiers of received packets. Consecutive packets in a sequence have different identifier values, and the identifier values of a sequence of packets are preferably an incremental sequence. The packet identifier is preferably a number, but it may alternatively be some other symbol. The sequence of identifier values should be known by the output of the link so that the output can correctly identify missing packets which have been suppressed by the input.

5 The packet identifiers of a sequence of packets may be a recurring sequence of N symbols. In this case the input transmits at least one packet in every N packets such that the output can determine the exact number of missing packets. This ensures bit count integrity between the input to the link and the output of the link. One example of a protocol in which the packet identifiers of a sequence of packets are a recurring sequence of N symbols is the ATM AAL1 protocol. In the ATM AAL1 protocol the identifier values of a sequence of consecutive packets are the recurring set of sequence numbers 0,1,2,3,4,5,6,7.

10 Preferably data for the missing packets is generated according to data carried in other packets received at the output. The data for the missing packets is generated according to data carried in received packets preceding the missing packets, and in its simplest form replicates data carried in received packets preceding the missing packets.

20 Preferably data is generated for a missing packet only if the missing packet follows one or more packets whose pattern of data indicates that the missing packet may not have been transmitted across the link. This has the advantage of avoiding errors caused by regenerating a packet which was not purposely suppressed at input, but lost by the transmission network. The packet preceding the missing packet may have a payload of data comprising a plurality of identical units, such as a sequence of idle flags; or a repeating pattern of n data bytes. Alternatively, the missing packet may be preceded by two cells having identical payloads.

25 The output may store data patterns such as commonly occurring flags or byte patterns, for use in matching with received data. Updated pattern information, for use in pattern matching, may be sent to the output by the input or by a network management unit via a management communications channel.

30 Another aspect of the invention provides a method of operating an output of a telecommunications link on which bandwidth is being minimised by not transmitting packets containing redundant data comprising the steps of:

- 35
- receiving a plurality of data packets, each packet comprising data and a packet identifier;

- determining missing packets according to the identifiers of received packets; and,

- generating data for the missing packets, which generated data corresponds to the redundant data not transmitted across the link.

5

A further aspect of the invention provides apparatus for minimising bandwidth used on a telecommunications link comprising, at an input to the link:

- means for receiving a plurality of data packets, each packet comprising data and a packet identifier;

10

- means for determining packets which contain redundant data;
- means for transmitting, across the link, only the packets which do not contain redundant data;

and further comprising, at an output of the communications link:

15

- means for receiving the transmitted packets;
- means for determining missing packets according to the identifiers of received packets; and,
- means for generating data for the missing packets, which generated data corresponds to the redundant data not transmitted across the link.

20

A still further aspect of the invention provides apparatus for use at an output of a telecommunications link on which bandwidth is being minimised by not transmitting packets containing redundant data comprising:

- means for receiving a plurality of data packets, each packet comprising data and a packet identifier;

25

- means for determining missing packets according to the identifiers of received packets; and,

- means for generating data for the missing packets, which generated data corresponds to the redundant data not transmitted across the link.

30

The following detailed description describes the bandwidth-saving feature with reference to an ATM network carrying constant bit rate services. However, the invention has more general application to other packet, cell or frame-based transport mechanisms such as Switched Multimegabit Data Service (SMDS) or Connectionless Broadband Data Service (CBDS). While there is more opportunity for bandwidth-saving with packets carrying constant bit rate traffic, the bandwidth-saving feature is not limited to being

35

used with packets carrying constant bit rate traffic, but can also be used with packets carrying variable bit rate traffic.

5 The packets, cells or frames may be full of user data or they may be only partially filled, such that part of a packet contains user data and the remainder of the packet contains padding data. Partial filling is a technique which is often used in delay-sensitive applications where it is undesirable to wait to fill an entire packet with user data.

10 Preferred features may be combined as appropriate, and may be combined with any of the aspects of the invention as would be apparent to a person skilled in the art.

15 **BRIEF DESCRIPTION OF THE DRAWINGS**

For a better understanding of the invention, and to show by way of example how the invention may be carried into effect, embodiments of the invention will now be described with reference to the accompanying drawings in which:

20 Figure 1 shows terminal equipments which are interconnected by a conventional time-division multiplexed network;

Figure 2 shows terminal equipments which are interconnected by an ATM network;

Figure 3 shows an example of a telecommunications network in which the invention may be used;

25 Figure 4 shows the interworking part (IWF) of figure 3 in more detail;

Figure 5 shows the format of cells which can be transmitted over the network of figure 3;

30 Figure 6 is an overview of the bandwidth-saving feature;

Figure 7 shows one example of the bandwidth-saving feature;

Figure 8 shows another example of the bandwidth-saving feature;

Figure 9 is a flow-chart to illustrate the suppression algorithm performed at the ingress of the ATM network;

35 Figure 10 shows one way of inspecting the payload of a cell;

Figure 11 is a flow-chart to illustrate the marking algorithm performed at the egress of the ATM network;

Figure 12 is a state machine for one way of performing the regeneration algorithm at the egress of the ATM network.

5 **DESCRIPTION OF PREFERRED EMBODIMENTS**

Many telecommunications customers require information to be transported at a well-defined, constant bit rate (CBR). Real-time voice is an example of a constant bit rate service. To meet this demand traditional circuit-switched networks have evolved by assigning bandwidth in multiples of 64Kbit/s regardless of whether this bandwidth is required for the entire duration of a call. Many such services in use today have been explicitly designed to use existing transport mechanisms which can transfer information in multiples of 64kbit/s. Such services are often called Nx64kbit/s services and usually comprise some form of time-division multiplexed (TDM) transmission network, as shown in figure 1.

However, as far as the customer is concerned, as long as information is provided at a constant rate at one end of the connection, and it arrives at the same rate at the other end of the connection, the transport mechanism across the network is of little or no concern. With the advent of ATM technology, there is now a demand for interoperability between ATM and customer's existing TDM equipment. To meet this demand it must be possible to carry the CBR over the ATM network, emulating the required circuit characteristics appropriately. Some form of conversion of the Nx64kbit/s information must take place before it can be transferred as ATM cells across an ATM network. This conversion is performed by an ATM Interworking Function (IWF) as shown in figure 2.

Figure 3 shows a communications network having a number of communications sources (video codec 30, PBX 34, Router 35) connected to similar devices 60, 61, 63 via a combination of a time-division multiplexed network and an ATM network 40, 45, 46.

A TDM Multiplexer 36 receives inputs from terminal equipment 31, PBX 34 and router 35. Each input comprises a data signal which is a multiple of 64kbit/s. Multiplexer 36 multiplexes the inputs into a time division multiplexed output data stream, such as a Plesiochronous Digital Hierarchy (PDH) DS1 or E1 link. An interworking unit IWF 38 receives the data from

the DS1/E1 links and packages it into a form for transmission over ATM network 40, 45, 46 to a further Inter-working unit IWF 48. IWF 48 converts the ATM packaged data to a form for onward transmission over a further DS1/E1 link to a further multiplexer 55. Mux 55 demultiplexes the data for transmitting to equipment 58, 61, 63.

As indicated on fig. 3, equipments 30-36 may be sited at a customer's premises, and the ATM network may be part of a public network.

Referring again to fig. 3, each input to Mux 36 must be a multiple of 64kbit/s. However, some sources generate a data stream which varies in bit rate. As an example, video codec 30 generates data at a variable rate according to the activity and detail in a video picture. Codec 30 connects to Mux 36 via a terminal equipment 31 which receives the variable bit rate data stream from codec 30 and generates a constant nx64bit/s output. TE 31 achieves this by inserting idle flags into the data stream during idle periods. In a protocol called High Data Link Control (HDLC) an idle flag is represented by the byte '01111110'. Idle periods may vary in length from being very short, such as where a video codec generates data at 30kbit/s but must be padded out to 64kbit/s, or very long, such as where a video conference has ended and codecs have been turned off but the connection has not been terminated.

Another condition which can occur is that a piece of terminal equipment, such as a telephone 33 generates an alarm signal, represented by a continuous stream of binary ones '1111.....'.

The operation of the TDM/ATM interworking units IWF will now be described in more detail, with reference to figure 4.

The interworking unit enables users of PDH terminal equipment and services to operate over an ATM network, obtaining a Quality of Service (QoS) comparable with that offered by PDH-based networks. This enables network operators to migrate to a broadband ATM network whilst continuing to offer support for existing PDH terminal equipment. PDH data is transported via the ATM Adaptation Layer (AAL) type-1. This service is tailored to continuous bit rate (CBR) services, known as class A, because it receives and delivers Protocol Data Units (PDUs) with the same constant bit rate to or from the next higher protocol layer.

Two transfer modes are available: unstructured mode and structured mode.

5 In unstructured mode an entire PDH bit stream, including the frame overhead bit(s), is conveyed point-to-point across the ATM network without interpretation using ATM Adaptation Layer type-1 (AAL1) cells sharing a single virtual channel across the ATM network. The complete PDH bit stream is inserted into the SAR-PDU payload of AAL1 cells without regard to ATM-PDH octet alignment.

10 In structured mode, the ATM network bandwidth demands can be reduced as the data from selected PDH time slots, corresponding to particular Nx64kb/sec services, is transported from the incoming PDH frames across the ATM network. Each Nx64kb/sec service uses a separate virtual channel across the ATM network. Data from particular PDH time slots is inserted into the SAR-PDU payload of AAL1 cells.

20 The interworking unit is shown with one DS1 or E1 Plesiochronous Digital Hierarchy (PDH) interface 105 coupled to a single ATM Adaptation Layer type-1 (AAL1) processing function within the AAL1 Processing module 130. The AAL1 Processing module is connected to separate ingress and egress ATM cell interfaces 140, 145.

25 PDH data received from each PDH interface is optionally framed and passed to the timeslot and queue function. In this function the incoming data bytes are assigned to channel queues according to their port and timeslot. If Channel Associated Signalling (CAS) mode is selected for a particular channel, signalling bytes are also assigned to the channel queues. The data and signalling are queued so as to generate the AAL1 data structure according to the mode of transmission selected. Once sufficient data as
30 been queued, the data structure is passed to the AAL1 processing function. The amount of data may be a full cell payload or a partial cell payload depending on whether partial filling is being used. When a PDH service interface is operating in unstructured mode, all the receive data is placed
35 directly in a single virtual channel queue.

The AAL1 processor 130 may add a pointer field to the data structure while generating AAL1 Segmentation and Reassembly Protocol Data Units (SAR-

PDUs) which are then placed within cells, together with control and optional timing information. A separate Virtual Path Identifier (VPI)/Virtual Channel Identifier (VCI) in the ATM cell header is used for cells of each service. The cells from each AAL1 processing function are then passed to an ingress ATM cell interface 140 where the ingress part of the bandwidth-saving feature is performed by a redundant data identification block and a cell suppression block. ATM cells are then queued for transmission over the ATM network.

At a destination, or egress ATM cell interface 145, a pattern identification block identifies cells containing data having a predetermined pattern, and marks these cells. The cells are then routed to an AAL1 processing function, via an internal cell demultiplexer, in accordance with the VCI/VPI labels contained in their ATM cell headers. The SAR-PDU is extracted and the recovered PDH data is placed within a 'play-back', or reassembly buffer which is read at a rate determined by either an external clock reference, or timing information recovered from the Protocol Control Information (PCI) held within the SAR-PDU header. A data regeneration block identifies cells which have been suppressed and regenerates data.

Figure 5 shows the structure of cells used in one format of AAL1 transmission. ATM cells are 53 bytes long with a 5 byte ATM cell header and a 48 byte payload. The ATM cell header CH includes information which determines the routing of the cell through the ATM network, i.e. the virtual channel. The payload portion of the AAL1 cell comprises 46 or 47 bytes of constant bit-rate data (depending on the AAL1 format used) and a 1 or 2 byte AAL1 header. The AAL1 header includes a sequence number SN and protection for the sequence number SNP. Within the sequence number block SN there is a sequence count SC which contains a value of 0, 1, 2, 3, 4, 5, 6 or 7. For a particular virtual channel, consecutive cells in time have consecutive sequence numbers. Consecutive cells have the sequence numbers 0, 1, 2, 3, 4, 5, 6, 7, 0, 1, 2.... i.e. the sequence number repeats after every 8 cells. The count of sequence numbers for each virtual channel across the ATM network is independent of the count used on other virtual channels.

Figure 6 gives an overview of the process for saving bandwidth in the ATM network. The effect of the bandwidth-saving process on a small group of cells is shown at various points in the system.

5 At the ingress, incoming data is packaged into cells by AAL1 processor 130, resulting in a stream of cells. Ingress ATM cell interface 140 identifies cells which contain redundant data, and suppresses them i.e. does not transmit them. This reduced group of cells is transmitted across the ATM network. At the egress of the ATM network the cells are received. ATM interface 145
10 identifies cells which contain data having a particular pattern, such as cells containing a string of idle flags. These cells are marked (M). Missing cells are identified by examining the sequence numbers of received cells. Cells which were suppressed by the ingress ATM cell interface 140 are regenerated by the egress AAL1 cell processor 130. Cells are only
15 regenerated when they are preceded by a marked cell, to avoid errors caused by regenerating data for cells which are missing due to reasons other than being suppressed at the ingress.

20 The egress IVF therefore outputs the same number of bits as were originally input at the ingress, these bits having exactly the same content as those originally input at the ingress. By suppressing cells at the ingress to the ATM network bandwidth is made available on the ATM network for other traffic.

Examples of redundant data are:

- 25
- cells containing the same pattern of data bits or bytes;
 - cells containing idle flags, where a source has become idle, but is transmitting idle flags to maintain a constant bit rate;
 - cells containing alarm flags, where a source has a fault and transmits a stream of alarm flags.

30

The suppression method works by identifying one or more cells which contain a payload of data that follows a predetermined or predictable pattern. The embodiment described here in detail identifies consecutive cells having identical payloads. The first of the cells which contain the same data is sent
35 and the subsequent matching cells are suppressed. At the egress, suppressed cells are identified by checking the sequence number of each received cell. A cell which has been suppressed at the ingress results in a missing cell sequence number in the stream of cells received at the egress.

Data for a suppressed cell is regenerated by copying the payload content of a received cell which preceded the missing cell.

5 Figures 7 and 8 show two examples of cell suppression and regeneration.
For clarity, these examples show a simplified case where there are only four
numbers in the list of cell sequence numbers. In figure 7, at the ingress the
cells having the sequence numbers 2 and 3 contain the same data. The cell
having sequence number 2 is transmitted and the cell having sequence
10 number 3 is suppressed, such that the cells having the sequence numbers 1,
2, 4, 1, 2... are transmitted by the ingress and received at the egress. At the
egress it is recognised that a cell having sequence number 3 is missing, and
that cell together with its payload of data is regenerated.

15 Figure 8 shows an example where each cell in a long stream of cells
contains the same pattern of data. The first cell, having sequence number 1,
is transmitted and the cell having sequence number 2 is suppressed. The
ingress cannot suppress all of the following cells (having sequence numbers
3, 4, 1, 2,...) which contain the same data pattern because the egress would
20 not be able to identify all of the missing cells. This would not ensure bit
count integrity across the network. Therefore, one cell in every four is sent
by the ingress (four being the number of elements in the list of sequence
numbers.) Thus, the ingress transmits cells having the sequence numbers 1,
1, 1, 1,... The egress can identify that three cells are missing before each
25 received cell having the sequence number 1 and correctly regenerate data to
ensure bit count integrity. It will be appreciated that this technique can be
applied more generally to systems where there are N numbers in the list of
sequence numbers, and at least one cell in every N cells are transmitted by
the ingress. In the AAL1 protocol N=8.

30 The above examples show how cells which carry the same payload of data
can be suppressed. In the more general case cells having payloads of data
which follow a predetermined, or predictable, pattern can be suppressed. To
achieve this the ingress looks for data following a predetermined pattern,
suppresses redundant cells and the egress regenerates data based on the
35 payload of the received cell preceding the missing cell and a generator which
operates using the same rules as used in the ingress pattern matching unit.

One method of performing suppression at the ingress is now described with reference to figure 9. This method identifies cells having identical payloads and where the payload of each cell comprises a set of identical bytes. Firstly, at step 305, a check is made as to whether the suppression feature is enabled. The suppression feature is provisionable on a per-channel basis. At step 310 the payload of a cell is examined to see if all of its bytes are identical. If a cell is only partially filled with user data, it is only the part of the cell carrying user data that is examined. If the bytes are all the same, the method proceeds to step 310 where a check is made as to whether all of the payload bytes of the previous cell were identical. A flag 'Previous Match' indicates whether this is so. If so, the method proceeds to step 330 where a check is made whether the payload of this cell is the same as the previous cell. This is achieved by comparing the Payload Pattern of this cell with the previous one. If this is true, then the current cell matches the previous one identically, and could be suppressed. The last check is to check whether seven consecutive cells have already been suppressed. If less than seven cells have been suppressed, the cell under test is suppressed and a match counter is incremented, to track the number of suppressed cells. The purpose of this last step is to ensure that one cell in every eight is transmitted so that the egress can identify the correct number of missing cells, as explained earlier with reference to figure 8. If the answer to tests 310, 320, 330 are false then the match counter is set to zero and the cell under test is sent.

The primary reason for a source to emit a continuous stream of identical data is an idle condition. In this condition, each byte of a cell will be the same. This may be the HDLC idle flag '01111110' or some rotated version of this flag such as '11001111'. Rotation is caused by incoming data bytes not being aligned with the bytes of the AAL1 cell payload. Figure 10 shows the process for checking whether all of the bytes in a cell's payload are the same, and equal to a predetermined byte such as an idle flag. The first byte D1 of the payload after the sequence number SN is compared with the idle flag byte in each of its rotated orientations until a match is found. All other payload bytes (D2, D3.... D47) are also compared with the provisioned byte in the same orientation.

The method just described looks for cells which contain an identical set of data bytes. The byte could be a commonly occurring byte such as an idle

flag which is stored at the ingress for ready comparison with the data, or any other byte which occurs.

5 In order to maximise bandwidth savings, the ingress and egress should be able to identify a wide range of patterns representative of redundant data in the data carried in cell payloads. Some other patterns will now be described.

- repeating pattern of n bytes within a cell.

10 The method described above for checking a single byte can be extended to look for repeating patterns of n bytes in length. A set of commonly occurring patterns can be stored in a look-up table at the ingress for ready comparison with the cell data, or the ingress can examine the cell data to look for any repeating data pattern. The egress pattern identification block similarly examines the data in received cells to identify repeating n byte patterns so
15 as to mark the cell.

- repeating cell payload

Consecutive cell payloads (not necessarily made up of a set of identical data bytes) can be examined to look for a match. There are two possibilities:

20

- (i) a sequence of consecutive cell payloads match a predetermined stored payload pattern.

25 The ingress has a look-up table of data payloads that are compared with the incoming data. If a match is found then the ingress sends the first cell and suppresses the subsequent ones. At the egress the received cell payloads are compared with a similar look-up table. When a cell payload matches one of the stored payloads it is marked and if it is followed by a missing cell, the egress knows to replicate the cell payload.

30

- (ii) a sequence of consecutive payloads are identified, but do not match a predetermined stored pattern.

35 The ingress identifies that there are identical cell payloads and sends the first two cells and suppresses subsequent ones in the sequence. The egress receives two consecutive identical cell payloads followed by one or more missing cells. The egress is programmed to recognise this, and replicates the last received

payload. If only one cell were sent, the egress would not be able to identify that there was a pattern in the data and would therefore not regenerate any data for the missing cells.

5 • pattern spanning cell payloads.

For example, two consecutive cells at the ingress contain data which uniformly increases in magnitude (i.e. a data ramp), and the ingress is set to recognise this pattern as redundant data. The first cell is transmitted and the second cell is suppressed. The egress is also set to identify cells which
10 contain data that uniformly increases in magnitude. Upon receiving a cell in which the payload uniformly increases in magnitude followed by one or more missing cells the egress is able to regenerate the missing data by applying the rule of calculating data for the missing cells that uniformly increases in value, continuing on from the data contained in the last received cell.

15

At the egress, the process of regenerating cells comprises the processes of:

- inspecting the data payload of a cell to check whether it contains a pattern of data (e.g. a payload of identical bytes) which indicates that it may be followed by a missing cell; and,
- 20 - looking for missing cell sequence numbers.

The reason for performing both of these processes rather than relying just on identifying missing cell sequence numbers is that cells could be missing for reasons other than cells being purposely suppressed at the ingress; e.g.
25 cells may be misrouted by the ATM network or lost due to faults in the ATM network.

A missing sequence number preceded by a cell having a payload of identical bytes indicates that a cell has been suppressed with a high degree of
30 reliability.

The method of inspecting a cell payload is shown in fig. 11. Firstly, at step 400 a decision is made as to whether the suppression/regeneration feature is enabled. If so, at step 410 the payload of the received cell is examined to
35 see if it contains a pattern of data indicating that it may be followed by a missing cell. This can be a test which looks for a payload containing: a set of identical bytes; a repeating pattern of n bytes; a data ramp, or some other

data pattern. This step may refer to a look-up table containing commonly-occurring bytes or data patterns and compare these with the payload data.

5 When a cell is matched, the data regeneration block in AAL1 processor 130 is informed of the data pattern that was matched to allow it to accurately regenerate the data. A cell having a payload of identical bytes is marked.

10 One way in which a cell can be identified as 'marked' is to use a flag at a predetermined bit position in one of the bytes of the ATM cell header.

15 The egress pattern identification block 145 may receive information about new data patterns from the ingress or from a network management function using proprietary management communication channels. For example, where the ingress is aware that it is sending, or will be sending, repeated identical payloads of a particular data pattern it can inform the egress of this pattern, thereby allowing the egress to identify this pattern.

20 A state machine for processing received cells will now be described with reference to figure 12. This state machine is based on the Fast SN algorithm for sequence number processing described in ITU I.363.1 Appendix III. The state machine described in this appendix describes how to deal with events where a cell is lost or misinserted, and in the case where a cell is missing inserts a dummy cell.

25 A decision in this algorithm is taken immediately after the analysis of the received cell. This means that when a cell is received the SN is immediately evaluated and the cell is eventually passed to the final destination to minimise transmission delay. In the state machine the action always refers to the last cell received.

30 A valid SN is defined as an SN which has no detected errors or it had an error that was corrected. A Match cell is a cell which has been marked as a match by the ATM interface logic.

35 When dummy cells are inserted, the payload is the provisioned dummy data byte. When suppressed cells are regenerated the payload is the same as the previously received matched cell.

The following acronyms are used in the algorithm:

SN - Sequence Number Including error protection

SC - Sequence Count value

5 When inserting dummy cells or regenerating suppressed cells the number of cells to be inserted can be calculated from the gap in the sequence number series. For example, receiving cells with the sequence numbers:

1 2 3 7 0

10 ideally requires 3 cells to be inserted in between the cells having sequence numbers 3 and 7.

The details of the algorithm are the following:

a) START

15 It is the initial state. It remains in this state discarding cells until there is a valid SN.

- if the cell is Match then goto OUT OF SYNC (Match);
- if the SN is invalid remain in this state;
- otherwise goto OUT OF SYNC (No Match);

20

b) OUT OF SYNC (No Match)

In this state the sequence counting is not synchronized yet.

- if the SC is in sequence with the previous one, it goes to SYNC (No Match) or SYNC (Match) as appropriate; the received cell is accepted;
- 25 • if the SC is out of sequence and not a match then remain in this state, discarding the received cell;
- if the SC is out of sequence and a match goto OUT OF SYNC (Match), discarding the received cell;
- if the SN is invalid return to START, discarding the received cell;

30

c) OUT OF SYNC (Match)

Again sequence counting is not yet fully synchronized, but because we have had a matched cell we must assume that if an out of sequence cell arrives it is because cells were suppressed at the ingress to the network.

- 35 • if the SC out of sequence and a match then remain in this state, regenerate suppressed cells and accept received cell;
- if the SC out of sequence and not a match then goto OUT OF SYNC (No Match), regenerate suppressed cells and accept received cell;

- if the SC is in sequence with the previous one then goto SYNC (Match) or SYNC (No Match) as appropriate, accepting the received cell;
- If the SN is invalid then goto START, discarding the received cell;

5 d) SYNC (No Match)

In this state the sequence counting is considered to be synchronized:

- if the SC is in sequence with the previous one and the cell is not a match then remain in this state accepting the received cell;
- if the SC is in sequence with the previous one and the cell is a match,
10 goto SYNC (Match) and accept the received cell;
- if the SC is not in sequence with the previous one and the cell is not a match, goto OUT OF SEQUENCE (No Match) accepting the received cell;
- If the SC is not in sequence with the previous one and the cell is a match, goto SYNC (Match) inserting dummy cells and accepting the received
15 cell;
- if the SN is invalid then goto INVALID (Match) or INVALID (No Match) as appropriate, accepting the received cell;

e) SYNC (Match)

20 In this state the sequence counting is considered to be synchronised but we assume that any out of sequence cells are caused by suppression at the ingress to the network.

- if the SC is in sequence with the previous one and the cell is not a match, goto SYNC (No Match) accepting the received cell;
- 25 • if the SC is in sequence with the previous one and the cell is a match, remain in this state accepting the received cell;
- if the SC is out of sequence with the previous one and the cell is not a match, goto SYNC (No Match) regenerating the suppressed cells and accepting the received cell;
- 30 • if the SC is out of sequence with the previous one and the cell is a match, remain in this state regenerating the suppressed cells and accepting the received cell;
- if the SN is invalid goto INVALID (Match) or INVALID (No Match) as appropriate, accepting the received cell; (if the cell was a match cells might
35 have been suppressed but because of the invalid SN it is impossible to know)

f) OUT OF SEQUENCE (No Match)

In this state the following actions are taken when a cell arrives:
NB. there is no OUT OF SEQUENCE (Match) state.

- if the SN is invalid, discard the received cell and goto START;
- 5 • if the SN is valid and the SC is in sequence with the last cell received with a valid SN, goto SYNC (Match) or SYNC (No Match) as appropriate; the received cell is considered to be misinserted and is discarded to keep bit count integrity because the previous cell is considered to be misinserted but is already sent;
- 10 • if the SN is valid and the SC is in sequence with the SC of the previous cell, the system assumes that cells were lost; insert dummy cells, accept the received cell and goto SYNC (Match) or SYNC (No Match) as appropriate;
- 15 • if the SC is valid and the SC has a value exceeding by two the SC value of the last cell received in sequence (i.e. the SN error protection mechanism has failed) assume that the received cell was in sequence and therefore accept the received cell and goto SYNC (Match) or SYNC (No Match) as appropriate;
- 20 • if the SN is valid but is not in any of the two previous situations, discard the received cell and goto OUT OF SYNC (Match) or OUT OF SYNC (No Match) as appropriate;

g) INVALID (No Match)

In this state the system shall take the following decisions on the received cell:

- 25 • if the SN is again invalid, goto START and discard the received cell;
- if the SN is valid and the SC is in sequence with the last cell received with a valid SN, goto SYNC (Match) or SYNC (No Match) as appropriate, but discard the received cell to keep bit count integrity because the previous cell
- 30 is considered to be misinserted but is already sent;
- if the SN is valid and the SC has a value exceeding by two the SC of the last cell received with a valid SN, it is assumed that although there was an invalid SN the received cell is in sequence so therefore accept it and goto SYNC (Match) or SYNC (No Match) as appropriate;
- 35 • if the SN is valid but is not in any of the previous situations, discard the received cell and goto OUT OF SYNC (Match) or OUT OF SYNC (No Match) as appropriate;

h) INVALID (Match)

In this state it is assumed that the received cell had a SC which was one more than the previously received cell with a valid SN;

- if the SN is again invalid, goto START and discard the received cell;
- if the SN is valid, goto SYNC (Match) or SYNC (No Match) as appropriate, inserting regenerated cells if necessary and accepting the received cell.

CLAIMS

1. A method of reducing bandwidth used on a telecommunications link comprising, at an input to the link, the steps of:
- 5 - receiving a plurality of data packets, each packet comprising data and a packet identifier;
- determining packets which contain redundant data;
- transmitting, across the link, packets which do not contain redundant data;
- 10 and further comprising, at an output of the telecommunications link, the steps of:
- receiving the transmitted packets;
- determining missing packets according to the identifiers of received packets; and,
- 15 - generating data for the missing packets, which generated data corresponds to redundant data not transmitted across the link.
2. A method according to claim 1 wherein the step of determining packets which contain redundant data comprises identifying a sequence of consecutive packets containing data that follows a predictable pattern and transmitting only the first packet or packets in the sequence.
- 20 3. A method according to claim 1 or 2 wherein identifying a sequence of consecutive packets containing data that follows a predictable pattern comprises identifying a sequence of packets containing same data.
- 25 4. A method according to claim 3 further comprising identifying a packet whose data comprises a plurality of identical units.
- 30 5. A method according to claim 4 wherein each unit comprises one of: an idle data flag; an alarm flag.
- 35 6. A method according to claim 4 wherein each unit comprises a rotated version of one of: an idle data flag; an alarm flag.
7. A method according to any preceding claim wherein only a portion of each packet contains user data and the step of determining packets which

contain redundant data examines only the portion of the packet containing user data.

5 8. A method according to any preceding claim wherein the packet identifiers of a sequence of packets are a sequence of symbols and the step of determining missing packets according to the identifiers of received packets comprises identifying missing symbols in the sequence of identifiers of received packets.

10 9. A method according to claim 8 wherein the packet identifiers of a sequence of packets are a recurring sequence of N symbols, and the input transmits at least one packet in every N packets.

15 10. A method according to claim 9 wherein the packets are ATM AAL1 cells and the packet identifiers are a sequence of 8 numbers.

20 11. A method according to any preceding claim wherein data for the missing packets is generated according to data carried in other packets received at the output.

 12. A method according to claim 11 wherein data for the missing packets is generated according to data carried in received packets preceding the missing packets.

25 13. A method according to claim 12 wherein the input identifies a sequence of packets containing same data and only transmits the first packet in the sequence and the output generates data for the missing packets by replicating data carried in received packets preceding the missing packets.

30 14. A method according to any preceding claim further comprising generating data for a missing packet only if the missing packet follows one or more packets whose pattern of data indicates that the missing packet may not have been transmitted across the link.

35 15. A method according to claim 14 comprising generating data for a missing packet if the missing packet follows a packet whose pattern of data matches a predetermined pattern stored at the egress.

16. A method according to claim 15 further comprising sending pattern information to the egress for use in pattern matching.

5 17. A method according to claim 14 comprising generating data for a missing packet only if the missing packet follows a packet whose data comprises a plurality of identical units.

10 18. A method according to any preceding claim wherein the packets are ATM cells.

19. A method of operating an output of a telecommunications link on which bandwidth is being reduced by not transmitting packets containing redundant data, comprising the steps of:

- 15 - receiving a plurality of data packets, each packet comprising data and a packet identifier;
- determining missing packets according to the identifiers of received packets; and,
- generating data for the missing packets, which generated data corresponds to redundant data not transmitted across the link.

20 20. Apparatus for use at an output of a telecommunications link on which bandwidth is being reduced by not transmitting packets containing redundant data comprising:

- 25 - means for receiving a plurality of data packets, each packet comprising data and a packet identifier;
- means for determining missing packets according to the identifiers of received packets; and,
- means for generating data for the missing packets, which generated data corresponds to redundant data not transmitted across the link.

30 21. Apparatus for reducing bandwidth used on a telecommunications link comprising, at an input to the link:

- 35 - means for receiving a plurality of data packets, each packet comprising data and a packet identifier;
- means for determining packets which contain redundant data;
- means for transmitting, across the link, packets which do not contain redundant data;

and further comprising, at an output of the communications link:

- means for receiving the transmitted packets;
- means for determining missing packets according to the identifiers of received packets; and,

5


- means for generating data for the missing packets, which generated data corresponds to redundant data not transmitted across the link.

Fig.1.



Fig.2.





3/10

Fig.4.

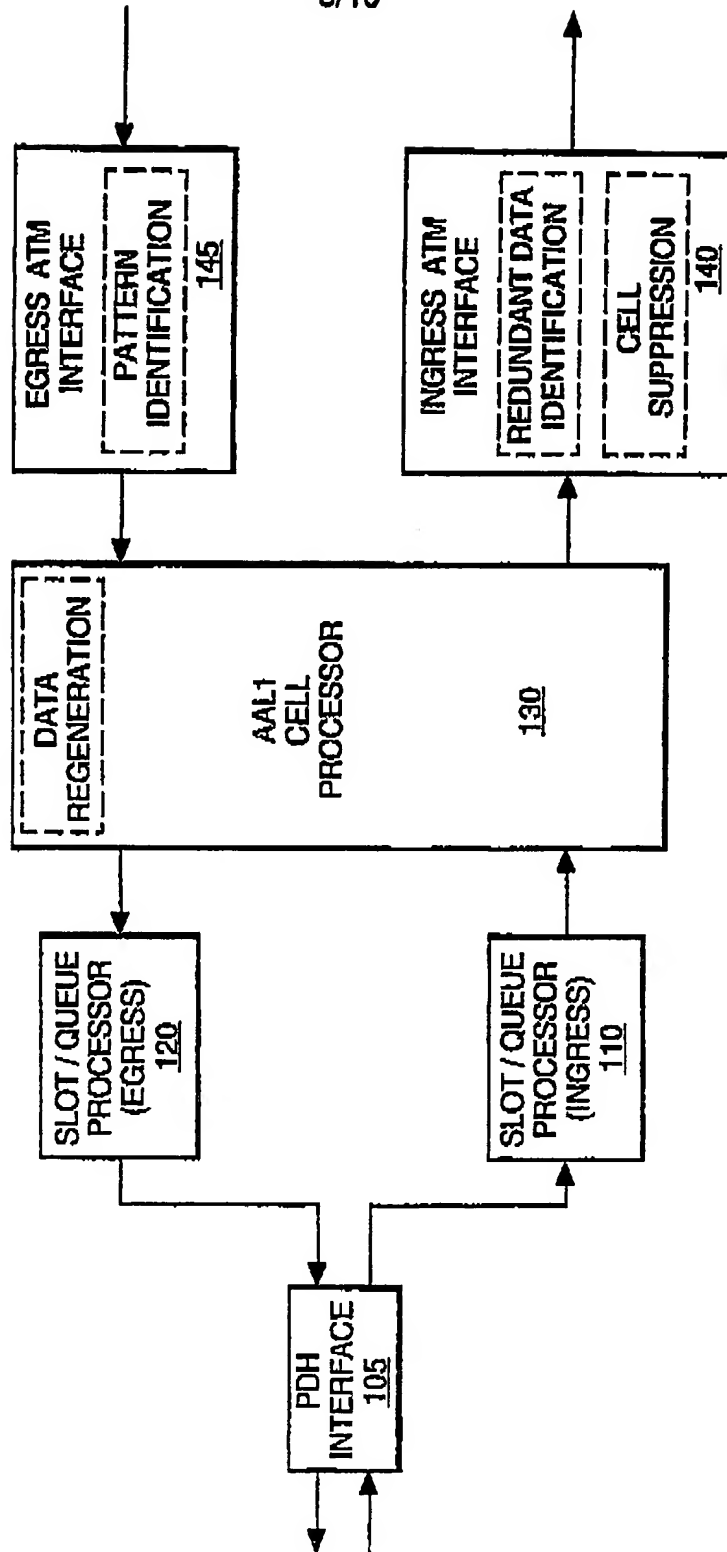


Fig.5.




Fig.6.




Fig.7.

SAME PAYLOAD
OF DATA

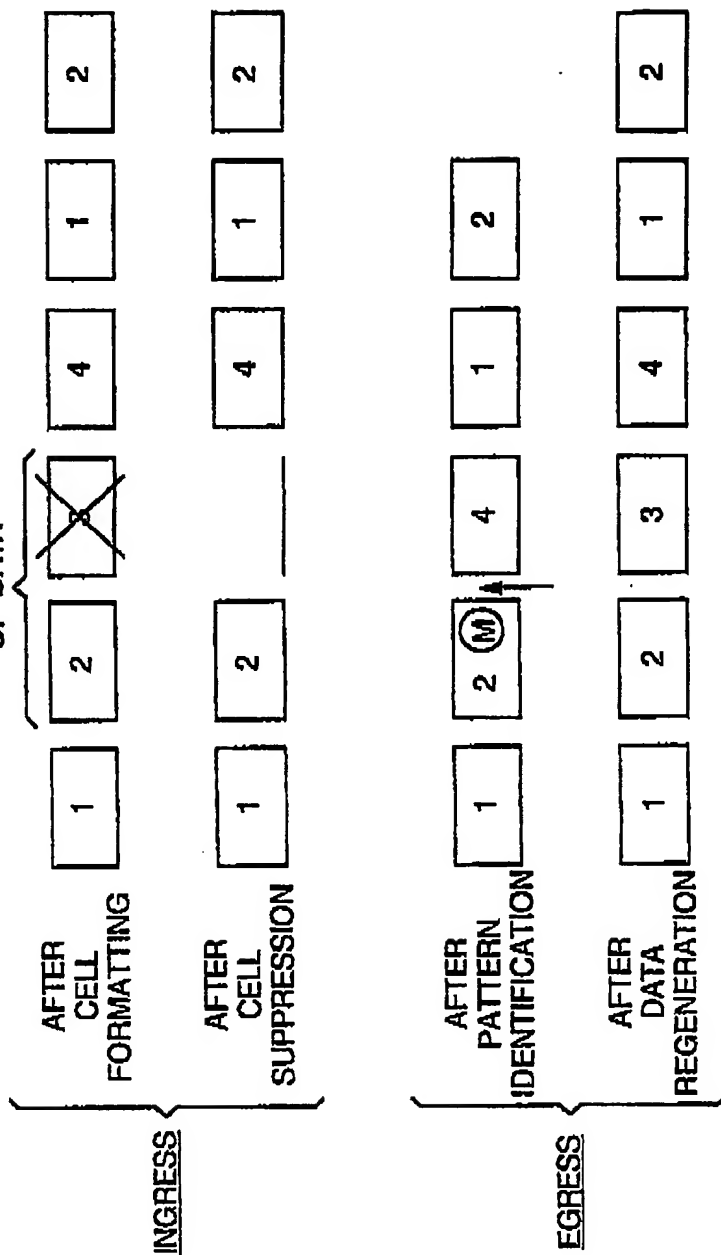

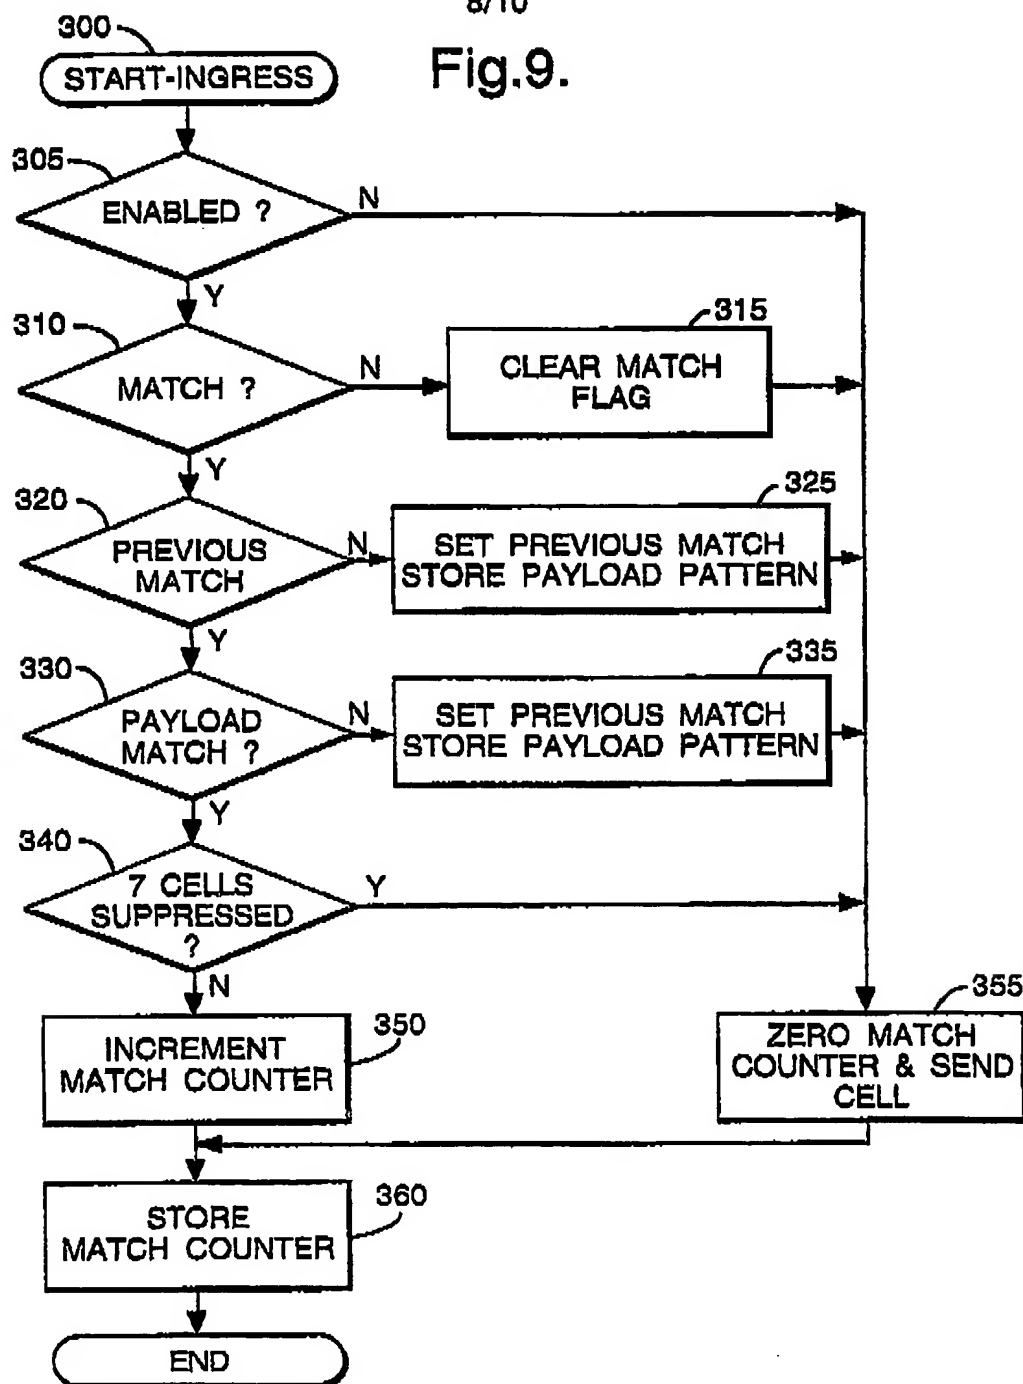


Fig.8.



8/10

Fig.9.



MATCH=ALL BYTES IN PAYLOAD ARE IDENTICAL
 PREVIOUS MATCH=PREVIOUS CELL HAD ALL BYTES IN
 PAYLOAD IDENTICAL
 PAYLOAD MATCH=NEW CELLS PAYLOAD MATCHES THE
 PAYLOAD OF THE PREVIOUS CELL

Fig.11.




Fig.10.




Fig.12:

